Monte Carlo Connection Prover

Michael Färber Cezary Kaliszyk
Universität Innsbruck
{michael.faerber,cezary.kaliszyk}@uibk.ac.at

Josef Urban
Czech Technical University in Prague
josef.urban@gmail.com

arXiv:1611.05990v1 [cs.LO] 18 Nov 2016

Abstract

Monte Carlo Tree Search (MCTS) is a technique to guide search in a large decision space by taking random samples and evaluating their outcome. In this work, we study MCTS methods in the context of the connection calculus and implement them on top of the leanCoP prover. This includes proposing useful proof-state evaluation heuristics that are learned from previous proofs, and proposing and automatically improving suitable MCTS strategies in this context. The system is trained and evaluated on a large suite of related problems coming from the Mizar proof assistant, showing that it is capable to find new and different proofs. To our knowledge, this is the first time MCTS has been applied to theorem proving.

1. Introduction

Proof automation is critical in formal proof verification. In particular, the recent *hammer* integration of automated theorem provers (ATPs) such as E (Schulz 2013), Vampire (Kovács and Voronkov 2013), CVC4 (Barrett et al. 2011), leanCoP (Otten 2008), and interactive theorem provers (ITPs) such as Isabelle (Wenzel et al. 2008), HOL4 (Slind and Norrish 2008), HOL Light (Harrison 1996), and Mizar (Grabowski et al. 2010) have led to a significant improvement of the usability of ITPs (Blanchette et al. 2016a, Blanchette et al. (2016b), Gauthier and Kaliszyk (2015), Kaliszyk and Urban (2014), Kaliszyk and Urban (2015a)). Today's automated theorem provers are however still quite weak in finding more complicated proofs over large formal developments. Their search typically blows up after several seconds, making the chance of finding proofs in longer times exponentially decreasing.

[Copyright notice will appear here once 'preprint' option is removed.]

This behaviour is reminiscent of poorly guided search in games such as Chess and Go. The number of all possible variants there typically also grows exponentially, and intelligent guiding methods are needed to focus on exploring the most promising moves and positions. The guiding method that has recently very significantly improved computer Go is Monte Carlo Tree Search (MCTS), i.e., expanding the search based on its (variously guided) random sampling.

In this work, we study MCTS methods that can guide the search in automated theorem provers, and evaluate their impact on interactive theorem proving problems. We focus on the connection calculus where the notion of tree search occurs naturally, and on the leanCoP prover, which has a compact implementation that is easy to experiment with. We can also build on previous machine-learning extensions of leanCoP (Kaliszyk and Urban 2015b, Urban et al. (2011)). As shown for example in the recent AlphaGo system (Silver et al. 2016), machine learning can be used to train good position evaluation heuristics even in very complicated domains that were previously thought to be solely in the realm of "human intuition". This gives a hope for productive use of MCTS in theorem proving. While "finishing the randomly sampled game" – as used in the most straightforward MCTS for games – is rarely possible in ATP (it would mean finishing the proof), there is a chance of learning good *proof state evaluation heuristics* that will guide the MCTS for ATPs in a similar way as e.g. in AlphaGo.

The rest of the paper is organized as follows. Section 2 introduces MCTS and in particular its recent UCT refinement. Section 3 introduces the connection calculus and proposes modifications that allow combining it with tree search. Section 4 proposes proof state evaluation heuristics. Section 5 describes our implementation, and section 6 evaluates the system on a set of Mizar Mathematical library problems.

2. Monte Carlo Tree Search

Monte Carlo Tree Search (MCTS) is a procedure to search large spaces by random sampling, biased towards promising parts of the search space. It has been applied

to a multitude of games and domains, often achieving state-of-the-art results (Browne et al. 2012). One of the most famous MCTS applications is AlphaGo, which is the first program to beat a professional Go player in a tournament without handicap (Silver et al. 2016).

The idea of MCTS is to run random play-outs of a game, called *simulations*, and to choose the move from where the play-outs were most successful. To obtain good results, it is crucial to bias the simulations such that they cover a diverse selection of the state space (*exploration*), but also that they concentrate on the parts where the most successful play-outs happened (*exploitation*). One of the most influential methods to balance exploration and exploitation in MCTS has been UCT (Kocsis and Szepesvári 2006). It contributes a strategy to select a subset of the decision space to search next, based on how often the subspace was already searched and what the average outcome of the searches was.

We next introduce terminology to describe MCTS problems and present MCTS and UCT algorithms in detail.

2.1 Characterisation of MCTS Problems

To make a problem such as theorem proving tractable by MCTS, it is convenient to characterise the problem by *rules* and *heuristics*. Rules describe the starting state and the legal moves possible in any given state, while heuristics estimate the quality of a given move as well as of a (final) state.

To describe rules, we use:

- A set of states $\mathcal{S}$.
- An initial state $s_0 \in \mathcal{S}$.
- A state transition function $\delta : \mathcal{S} \to 2^{\mathcal{S}}$, which maps every state to its permitted successor states.

To describe heuristics:

- A transition weight function $\delta_w : \mathcal{S} \to \mathcal{S} \to \mathbb{R}_{\geq 0}$, where $\delta_w(s,t)$ gives the weight of going from state s to t in a random simulation.
- A reward function $r : \mathcal{S} \to [0,1]$.

2.1.1 Example: Travelling Salesman Problem

As an example of an MCTS application, consider the travelling salesman problem (TSP): Given a list of cities $\{c_1, \ldots, c_n\}$ with a distance function d, find a permutation X of $[1, \ldots, n]$ such that $r(X) = \sum_i d(c_{X_i}, c_{X_{i+1}})$ is minimal. We can formulate the rules of this problem as follows:

- The set of states $\mathcal{S}$ are all possible (potentially unfinished) paths of the travelling salesman, i.e. permutations of subsets of $\{1, \ldots, n\}$.
- The initial state s_0 is the empty sequence $[]$.
- The state transition function δ returns all possible paths where the travelling salesman has visited one more city, i.e. $\delta([X_1, \ldots, X_i]) = \{[X_1, \ldots, X_i, j] \mid 1 \leq j \leq n \wedge j \notin \{X_1, \ldots, X_i\}\}$.

The heuristics for the TSP could be set as follows:

- The reward function is $r(X)$ as given in the TSP definition.
- The transition weight function is

$$\delta_w(s,t) = |\{u \in \delta(s) \mid r(u) \geq r(t)\}|^{-1},$$

which rates cities by how many cities are nearer to the current position of the travelling salesman.

This description is sufficient to run MCTS to find solutions to the TSP.

2.2 Monte Carlo Tree Search Algorithm

MCTS maintains a decision tree in which nodes correspond to states $\mathcal{S}$ of the underlying problem (say, the current positions of stones on a Go board), and the children of a node s correspond to successor states $\delta(s)$ (e.g., the board configurations after the other player has set a stone). In addition to the state, a node contains statistics about the number of times the node was traversed and about the rewards of simulations that were started from and below the node.

When starting MCTS, the decision tree is initialised with a root node that corresponds to the initial state $s_0 \in \mathcal{S}$. One iteration of the MCTS algorithm involves the following steps:

1. **Selection**: The MCTS tree is traversed from the root until a leaf (with at least one unexplored successor state) is encountered.
2. **Simulation**: A (random) simulation is performed starting from the leaf.
3. **Expansion**: New nodes corresponding to the initial simulation steps are added to the leaf.
4. **Backpropagation**: The reward of the simulation is backpropagated from the leaf to all nodes up to the root.

The actual algorithms used in all four steps are referred to as *strategies*. Some commonly used strategies are:

- **Simulation strategy**: Given a state s, choose as next state some state $t \in \delta(s)$ randomly biased by $\delta_w(s,t)$. If $\delta(s)$ is empty or a maximal simulation depth has been reached, stop the simulation.
- **Expansion strategy**: Add a node corresponding to the first state of the simulation to the leaf.
- **Backpropagation strategy**: Add the reward of the simulation to the rewards of all ancestors and increase their visit counts by one. This allows calculat-

ing the average reward by dividing the reward sum by the number of visits.

We present one of the most influential selection strategies: UCT.

2.3 UCT

UCT (Upper Confidence Bounds applied to Trees) is a version of MCTS with a selection strategy that attempts to balance exploitation and exploration: It selects the node j that maximises $\overline{X_j} + C_p\sqrt{\frac{2\ln n}{n_j}}$, where:

- $\overline{X_j}$ is average reward for node j,
- n_j is number of times node j was chosen,
- n is total number of times the parent node of j was chosen.

UCT assumes that rewards are always in the range $[0, 1]$. This is important to keep the equilibrium between exploration and exploitation. We call $C_p\sqrt{\frac{2\ln n}{n_j}}$ the *exploration term*.

3. Connection Proving and MCTS

In this section, we introduce the clausal connection calculus, discuss possible ways to model the connection calculus as game and describe how to integrate it with MCTS.

3.1 Clausal Connection Calculus

The connection calculus is a goal-directed proof search procedure (Bibel 1993). It operates on a matrix M representing the set of clauses obtained from the Skolem normal form of a set of first-order formulas. A version that is used in the automated theorem prover leanCoP is shown in Figure 1. The proof search is initialised with the "Start" rule and is successful when all branches of the proof tree are closed by the application of the other rules. The "state" of a branch is characterised by a triple $(C, M, Path)$, where C is the current clause, M is the matrix containing all input clauses and $Path$ is the active path.

3.2 Connection Proving as a Game

MCTS has been applied to two-player games such as Go (Silver et al. 2016), as well as to single-player games such as SameGame (Schadd et al. 2012). It is possible to model proof search as game with either one or two players. We compare the two approaches.

In case of a single player, the player chooses an action to apply to an open proof branch, potentially adding new open proof branches. The player wins when proof branches have been closed, and loses if there exists a proof branch for which no proof rule is applicable.

In case of two players, the first player chooses an open proof branch and, if the clause of the branch is not empty, a literal. The second player then applies an action to the chosen proof branch, which has to result either in closing the proof branch or producing at least one new open proof branch with the literal removed from the clause. If the second player is unable to do either of the two, then the first player wins, and if there is no open proof branch left for the first player to choose from, then the second player wins.

Modelling connection proving as a two-player game has the advantage that given an intelligent proof branch evaluation heuristic, the proof search can be directed by the first player towards "harder" proof branches, thus more quickly resulting in the termination of proof search in branches that cannot be closed. On the other hand, a similar behaviour can be achieved by pre-ordering the proof branches in the single-player game. However, if we model connection proving as a two-player game following the default UCT selection policy, *every* possible move of the first player has to be considered at least once. Even if one restricts the first player to choose only literals from a single proof branch clause, this can considerably blow up the search space. We therefore focus on connection proving as a single-player game.

3.3 Single-Player Connection Proving Game

To model single-player connection proving as MCTS problem, we introduce a non-branching version of the connection calculus (see Figure 2) where substitutions are explicit and every rule has only a single premise. As closed proof trees in the branching and non-branching connection calculi are isomorphic, the non-branching calculus preserves soundness and completeness. With this calculus, we can give the rules of connection proving as in subsection 2.1:

- The set of states $\mathcal{S}$ is the set of tuples $(Goals, \sigma)$, where $Goals$ is a set of clause-path pairs and σ is a substitution.
- The initial state s_0 is $(\{(\{\top\}, \{\})\}, \{\})$. $\neg\top$ has to be added to all positive clauses in M to correctly emulate the start rule of the branching connection calculus.
- The state transition function $\delta(s)$ returns all possible premises of non-branching connection calculus rules that have s as conclusion.[1]

In the remainder of this and the following section, we discuss possible MCTS heuristics.

3.4 Main Loop

To find good moves for two-player games, MCTS is usually run with a fixed time limit or number of iterations in order to select and use the best first move from the

[1] In practice, $\delta(s)$ restricts to a finite subset of the possible calculus premises, just like leanCoP does.

$$\frac{}{\{\},\ M,\ Path}\quad \text{Axiom}$$

$$\frac{C,\ M,\ \{\}}{M}\quad \text{Start}\quad \text{where } C \in M, C \text{ is positive}$$

$$\frac{C,\ M,\ Path \cup \{L_2\}}{C \cup \{L_1\},\ M,\ Path \cup \{L_2\}}\quad \text{Reduction}\quad \text{where } \sigma(L_1) = \sigma(\overline{L_2})$$

$$\frac{C_2 \setminus \{L_2\},\ M,\ Path \cup \{L_1\} \qquad C,\ M,\ Path}{C \cup \{L_1\},\ M,\ Path}\quad \text{Extension}\quad \text{where } \begin{array}{c}\sigma(L_1) = \sigma(\overline{L_2}),\\ \sigma \text{ is rigid},\\ C_1 \in M, L_2 \in C_2,\\ C_2 \text{ is a copy of } C_1\\ \text{with variables renamed}\end{array}$$

Figure 1. The clause connection calculus used in leanCoP.

$$\frac{}{\{\},\ \sigma}\quad \text{End}$$

$$\frac{Goals,\ \sigma}{Goals \cup \{(\{\},\ Path)\},\ \sigma}\quad \text{Goal}$$

$$\frac{Goals \cup \{(C,\ Path \cup \{L_2\})\},\ \sigma_2}{Goals \cup \{(C \cup \{L_1\},\ Path \cup \{L_2\})\},\ \sigma_1}\quad \text{Reduction}\quad \text{where } \sigma_2(L_1) = \sigma_2(\overline{L_2}),\ \sigma_2 \text{ factors through } \sigma_1$$

$$\frac{Goals \cup \{(C, Path),\ (C_2 \setminus \{L_2\},\ Path \cup \{L_1\})\},\ \sigma_2}{Goals \cup \{(C \cup \{L_1\},\ Path)\},\ \sigma_1}\quad \text{Extension}\quad \text{where } \begin{array}{c}\sigma_2(L_1) = \sigma_2(\overline{L_2}),\ \sigma_2 \text{ factors through } \sigma_1,\\ C_1 \in M, L_2 \in C_2,\\ C_2 \text{ is a copy of } C_1\\ \text{with variables renamed}\end{array}$$

Figure 2. The non-branching clause connection calculus.

root. It is possible to use MCTS in such a way as an oracle in a conventional theorem prover: Given a prover state with high uncertainty as to which proof step is most appropriate (for example when having to select between several extension clauses), one can run MCTS and obtain the most promising next proof step, resuming the conventional proof search from there.

This approach has two problems. First, it is unclear in which situations it is beneficial to run MCTS instead of searching all proof options exhaustively. Second, extracting only the first best successor node of the Monte Carlo search tree discards all the statistics in the search tree about the quality of follow-up states.

For this reason, we follow a different approach, namely we use MCTS as *driving force* to control the proof search: In every MCTS iteration, a short proof search is performed, and the reward function controls which parts of the search space the proof search concentrates on, depending on the outcomes of the previous proof searches. A similar approach is taken in (Schadd et al. 2012).

As selection strategy, we use UCT as described in subsection 2.3, and as backpropagation strategy, we use the MCTS default strategy as described in section 2.

3.5 Simulation Strategies

Once the selection strategy has chosen a child node with at least one unexplored successor state, the simulation strategy chooses an unexplored successor state s and starts a random simulation. In every iteration of the simulation we calculate the set of successor states $\delta(s)$, i.e. all possible proof steps that can be applied in the current prover state s. Next, one of the successor states is randomly picked, biased with the transition weight δ_w, and the simulation is continued with the picked successor state. During the simulation *no* backtracking is done, as it would require storing search history information in the prover state.

The simulation returns the sequence of traversed states up to the current state s if either $\delta(s)$ is the empty set, i.e. no proof rule is applicable, or some maximum simulation depth is reached. In case that s has no more open subgoals, i.e. a proof has been found, the proof is returned.

To bias the selection of successor states, we assign a constant weight to reduction steps and focus on influencing the weight of extension steps in function of the corresponding extension clauses:

- Constant: Potential extension steps have the same weight.
- Inverse size: Extension clause is weighed with the inverse number of literals it contains.
- Rank: Extension clauses are weighed by the inverse of the number of smaller extension clauses.

Selection of extension steps using machine-learning, similar to FEMaLeCoP (Kaliszyk and Urban 2015b), is left as future work.

3.6 Expansion Strategies

Once the simulation strategy has finished with a resulting sequence of states, the expansion strategy is responsible for adding a new child node to the node n the simulation has started from. Since the chance of revisiting a node sharply decreases with its depth, most deeper nodes will not be visited again and storing them all takes up significant amount of memory. Therefore, we add only a single leaf node to the starting node n. Considering deeper subtrees in function of the expected number of future visits to the node, in order to minimise the costly recalculation of successor states, could be considered in future work.

Our leaf-adding expansion strategies are:

- First-node expansion: This strategy adds the first node of the simulation, i.e. the previously unvisited action of the chosen node to its parent node.
- Best-node expansion: This strategy takes the state of the simulation with the *lowest number of subgoals* and adds it to the chosen node. When a node has a lower number of subgoals than the starting node, at least one subgoal was closed. Adding only that node to the tree discards all other options to close the branch. This is similar to the *cut* heuristic in leanCoP, which makes proof search incomplete, but reduces the search space.

The first-node expansion strategy preserves completeness. Consider the leafs of a Monte Carlo search tree. They represent all possible ways to prove the original conjecture, and all of its successor states preserve this invariant. Since rewards are between 0 and 1, we can show that for $C_p > 0$, any leaf of the search tree will be selected at some point. This is because the exploration term of any node will become large enough for that node to be visited, regardless of its average reward.

4. Proof State Evaluation Heuristics

One of the important advantages of MCTS for games is that is does not require an evaluation function that estimates the win likelihood for a non-final state. For example, in a chess game it is not required to find a function that estimates how likely a player will win the game, but it is only necessary to determine at the end of a game which player won and perhaps by how much. However, in theorem proving, we are very rarely able to play a game to the end with a clear outcome, meaning that we still need to estimate the success of branches. Even if we encounter a branch where we are not able to continue to proof, thus being in a final state, that may not be considered a "loss" in the game sense because we might just have a single mistake that made us unable to finish the proof. It is clear that we need a more fine-grained notion of "success" to estimate such situations' potential to lead to a proof by choosing slightly different decisions.

4.1 Provability Estimates

We define the quality of a state as its *provability* estimate, i.e. the probability of finding a proof by pursuing that branch. We have devised several methods to estimate the quality of a branch. The linear combination of their estimates constitutes the reward function $r(s)$ of our Monte Carlo Connection Prover.

The first method is based on the fact that towards the end of a proof, the number of open subgoals goes to zero, and the number of open subgoals is always smaller or equal to the number of previously opened subgoals. So the first reward function for a state is just the ratio of open and previously opened subgoals.

The second method is based on data collected during previous proof attempts. We assume that certain goals are statistically easier to close than others, and also certain literals of goals are easier than others. This is similar to the approach taken in e.g. FEMaLeCoP (Kaliszyk and Urban 2015b), which prefers proof steps involving previously "successful" clauses. However, where FEMaLeCoP only needs to establish an order on a finite set of clauses that are applicable in a particular situation, we need to establish an order on a potentially infinite set of states. Furthermore, while FEMaLeCoP estimates the usefulness of a clause for the current state, we have to evaluate the state itself, putting it into relation with any other state. We describe the method in the next section.

4.2 Machine-learnt Provability Estimates

We say that a literal L_1 was proved successfully if either the reduction rule was applied to $C \cup \{L_1\}$ or the extension rule was applied and its left branch was completely closed (see Figure 1). A literal was proved unsuccessfully

when the reduction rule was not applicable and for all extension rule that were applicable, their left branches were not closed.

We now define our *literal provability* estimate: If we have not seen the literal l in previous proofs, then the provability of l is 1, that is, for unseen literals we are optimistic that we are able to prove them. If we have seen it in proofs being successfully proved for p and unsuccessfully proved for n times, then the success rate is $\frac{p}{p+n}$. To account for literals that we have rarely seen, we introduce the *certainty*

$$c(x) = 1 - C\frac{1}{x^D + 1}$$

(where C and D are constants) and calculate the total literal provability as

$$1 + c(p+n) \cdot \left(\frac{p}{p+n} - 1\right).$$

With the literal provability defined, we can define *clause provability*. A clause is proved if all of its literals are proved, therefore it is natural to define the clause provability as the product of its literals' provabilities.

We can apply the same idea to obtain the total prover state refutability – to solve the whole problem, it is necessary to prove all clauses of the prover state, therefore the total provability is the product of the clause provabilities. However, in our experiments it became quite clear that this very frequently gave rewards of 0, because a single literal in a single clause that had a bad ranking could deteriorate the reward. For that reason, we tried combining clause provabilities with the minimum function and with the harmonic, geometric, and arithmetic means.

5. Implementation

The implementation is based on the source code of the FEMaLeCoP prover (Kaliszyk and Urban 2015b), which is an OCaml version of leanCoP (Otten 2008) enhanced by machine learning techniques. For all the provers mentioned below, we implemented the usual leanCoP optimisations which are not part of the calculus, such as lemmas and the regularity check. Furthermore, all provers below use the same proof format as the leanCoP tactic described in (Kaliszyk et al. 2015) to automatically find proofs for HOL Light, therefore adapting the HOL Light proof tactic to use our Monte Carlo system should be straightforward.

5.1 Lazy List Connection Prover

The FEMaLeCoP implementation uses continuation-passing style (CPS) to implement backtracking supporting the Prolog "cut". While this is very efficient, it does not provide a simple representation of different ways to close a subgoal, because this information is "hidden" in the continuation. However, MCTS needs an explicit set of successor states to any prover state. To approach this, we first created a version of FEMaLeCoP that uses *lazy lists* instead of continuations to represent different proof options: To close a subgoal, there might be different reduction and extension steps. So for a given literal l, we return a lazy list of proof trees that prove l. We can then filter out only closed proof trees and backtrack by going through the list of proof trees. Cut can also be implemented by restricting the returned proof trees to the first closed one. We refer to the lazy list implementation as *lazyCoP*.

5.2 Non-branching Calculus

lazyCoP served as a base for an implementation of the non-branching calculus: In the non-branching calculus, we return as list of actions to close a subgoal only the next nodes of the proof subtrees and not the whole proof subtree to close the subgoal. This distinction is important, because it allows us to perform rapid random play-outs of the proof search without branching. The resulting implementation of non-branching proof search works also independently from MCTS, for example by using iterative deepening together with a depth-first search in the proof search decision tree.

5.3 Monte Carlo Connection Prover

Because the non-branching proof search corresponds already closely to the MCTS problem description in subsection 2.1, it is now relatively easy to replace the depth-first search tactic with MCTS. The largest adaptation necessary was the replacement of the global array-based substitution with a local substitution for every explored proof branch, where we chose a list-based substitution. We refer to the resulting system as *monteCoP*.

5.4 Monte Carlo Tree Search

We implemented the UCT variant of MCTS and validated its function with the travelling salesman problem as shown in subsubsection 2.1.1.

We added an optimisation in MCTS specific to automated theorem proving: When the list of unvisited and visited child nodes of a node is empty, the node itself is deleted. Without this deletion, nodes which have no hope of ever contributing to the proof search would still be visited from time to time, as the UCT exploration term of the nodes would grow.

Furthermore, we use the transition weight function δ_w not only to bias the selection of successor states *during* the simulation, but also when choosing an unexplored action of a Monte Carlo tree node.

5.5 Training Data

To collect literal refutation statistics as described in section 4, we created a lazyCoP version that reports statistics for the last iteration of the iterative deepening if a proof was found. To account for the same literal appearing among different clauses, we also register for each literal the clause it came from.

An important aspect is consistent Skolemisation, which ensures that constants introduced during Skolemisation bear the same name among different problems and prover runs. Initially, we stored the symbolic representations of clauses and literals in the literal statistics. However, the names of the Skolem constants can become so large that loading the training data took a considerable amount of time during proof search. We solved this problem by saving only hashes of the symbolic representations of literals and their clauses, which reduced the size of the training data by several magnitudes. To combine the training data from different examples, we add the positive and negative occurrences of each literal among all problems.

5.6 Related Work

randoCoP (Raths and Otten 2008) is another connection prover based on randomness. It runs leanCoP multiple times, shuffling the order of clauses and literals at the beginning of every proof search. The difference to our approach is that in randoCoP, the randomness only plays a role at the beginning of each proof search, after which the behaviour is deterministic. In monteCoP, however, the restart behaviour is intrinsically given by UCT, and randomness influences the proof search at every point where multiple options are available to close a subgoal.

6. Evaluation

We used the problems of the MZR@Turing division of CASC-J6 (Sutcliffe 2013). It consists of 1000 training problems that were known before the competition, as well as 400 testing problems that were used to evaluate the quality of the submitted competitors. All 1400 problems are taken from the MPTP2078 problem set (Alama et al. 2011), but sometimes augmented with facts that are not necessary for the proof to make proof attempts harder. We mapped the problems of the MZR@Turing division to their bushy counterparts (with minimised number of dependencies) of the MPTP2078 dataset and used these for evaluation.

6.1 Parameter Optimisation

First, we ran lazyCoP with timeout of 300s on the 1000 training problems, yielding 314 solved problems together with literal statistics as described in subsection 5.5. We then divided the solved problems into a parameter training set of 264 and a parameter testing set of 50 problems. On the parameter training set, we ran ParamILS (Hutter et al. 2009) with monteCoP, using a problem timeout of 2s. We ran 32 parallel ParamILS instances with different starting configurations, where each instance ran for a total time of 80000s. Next, we obtained the best resulting ParamILS configuration with respect to the number of inferences per solved problem on the parameter testing set.

6.2 Results

As baseline comparison, we chose a monteCoP configuration that resembles breadth-first search, by choosing a constant reward function and a simulation depth of 1. This causes the search tree to grow uniformly and encourages only exploration, as the reward of all nodes is the same. This version of monteCoP proves 41 problems of the benchmark. Next, we chose a monteCoP configuration with a higher simulation depth, still keeping the constant reward function and choosing potential next proof steps with equal probability. This corresponds to randomised beam search with increasing beam depth. In this configuration, monteCoP solves 49 problems. Finally, we guide the proof search by biasing extension clause selection towards smaller clauses, using the simulation strategy in subsection 3.5, and with the reward function introduced in section 4, using the literal provability data obtained from the lazyCoP runs on the MZR@Turing training problems. The resulting monteCoP configuration solves 64 problems, which is an increase of 30.6% compared to the unguided monteCoP proof search. The single best lazyCoP strategy solves 84 problems. Still, of the 64 problems proved by monteCoP, 15 solutions are unique, which is 17.9% of the problems that lazyCoP solves.

6.3 Inferences

We count the number of inferences as the number of successful extension steps. In monteCoP, we additionally count the number of MCTS iterations. This roughly corresponds to the number of nodes in the Monte Carlo search tree, i.e. the number of random proof searches, with the exception that our version of MCTS can also delete nodes during an iteration when they do not have any successor states.

The average number of monteCoP MCTS iterations is 5532. The average number of inferences is 320295 for lazyCoP and 117163 for monteCoP. On average, monteCoP performed 9.6 times as many inferences as lazyCoP on the problems that both provers solved. This can be explained by the fact that monteCoP always calculates all possible extension steps applicable, even if the extension step is not tried afterwards. Finding a way to compute fewer extension clauses might improve performance.

6.4 Experiments

In our experiments, the exploration constant C_p turned out to be very influential for proof search: Setting it high results in a strategy that resembles breadth-first search, while setting it low makes it resemble best-first search. Our hypothesis was that the optimal value of C_p is problem-dependent. For that reason, we experimented with an oscillating $C_p(i) = C_{p,0} + a\sin(2\pi\frac{i}{T})$, where i is the number of the current MCTS iteration, a is the oscillation amplitude and T is the oscillation period. Unfortunately, this did not yield a substantial improvement.

We were surprised to see that the best-node expansion strategy (see subsection 3.6) corresponding to the leanCoP *cut* strategy solved only 14 problems.

We tried random rewards in the range $[0, 1]$, which proved exactly as many problems as constant reward. This is because the selection strategy calculates the average reward from all previous simulations, which for the random rewards converges to 0.5 for all nodes.

(Schadd et al. 2012) suggests running MCTS multiple times with different random seeds to avoid MCTS getting stuck in local maxima. We tried this in combination with rerunning the same problem with multiple values of the exploration constant C_p for a shorter time, however, this did not achieve any improvement.

6.5 MPTP2078

In addition to the problems in the MZR@Turing dataset, we also ran an evaluation on all MPTP2078 problems. Here, the single best lazyCoP strategy solves 468 problems, whereas monteCoP solves 296. Of these 296, 56 are unique, i.e. 12.0% of the problems lazyCoP solves.

7. Conclusion

We described and implemented a connection prover based on Monte Carlo Tree Search with machine-learnt guidance heuristic. We evaluated it on two datasets, where on one dataset, our system was able to prove 17.9% new problems compared to a conventional connection prover. One can use our system as advisor in a theorem prover to choose between different possible proof steps. As future work, new prover state evaluation heuristics, e.g. using neural networks, might help to better estimate promising regions of the proof search.

References

J. Alama, D. Kühlwein, E. Tsivtsivadze, J. Urban, and T. Heskes. Premise selection for mathematics by corpus analysis and kernel methods. *CoRR*, abs/1108.3446, 2011. URL http://arxiv.org/abs/1108.3446.

C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In G. Gopalakrishnan and S. Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 171–177. Springer, 2011. ISBN 978-3-642-22109-5. doi: 10.1007/978-3-642-22110-1_14. URL http://dx.doi.org/10.1007/978-3-642-22110-1_14.

W. Bibel. *Deduction - automated logic*. Academic Press, 1993. ISBN 978-0-12-095835-1.

J. Blanchette, C. Kaliszyk, L. Paulson, and J. Urban. Hammering towards QED. *Journal of Formalized Reasoning*, 9(1):101–148, 2016a. ISSN 1972-5787. URL http://jfr.unibo.it/article/view/4593.

J. C. Blanchette, D. Greenaway, C. Kaliszyk, D. Kühlwein, and J. Urban. A learning-based fact selector for Isabelle/HOL. *J. Autom. Reasoning*, 57(3):219–244, 2016b. doi: 10.1007/s10817-016-9362-8. URL http://dx.doi.org/10.1007/s10817-016-9362-8.

C. Browne, E. J. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. P. Liebana, S. Samothrakis, and S. Colton. A survey of Monte Carlo tree search methods. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(1):1–43, 2012. doi: 10.1109/TCIAIG.2012.2186810. URL http://dx.doi.org/10.1109/TCIAIG.2012.2186810.

T. Gauthier and C. Kaliszyk. Premise selection and external provers for HOL4. In *Certified Programs and Proofs (CPP'15)*, LNCS. Springer, 2015. doi: 10.1145/2676724.2693173. URL http://dx.doi.org/10.1145/2676724.2693173. http://dx.doi.org/10.1145/2676724.2693173.

A. Grabowski, A. Korniłowicz, and A. Naumowicz. Mizar in a nutshell. *J. Formalized Reasoning*, 3(2):153–245, 2010.

J. Harrison. HOL Light: A tutorial introduction. In M. K. Srivas and A. J. Camilleri, editors, *FMCAD*, volume 1166 of *LNCS*, pages 265–269. Springer, 1996. ISBN 3-540-61937-2.

F. Hutter, H. H. Hoos, K. Leyton-Brown, and T. Stützle. ParamILS: an automatic algorithm configuration framework. *Journal of Artificial Intelligence Research*, 36:267–306, October 2009.

C. Kaliszyk and J. Urban. Learning-assisted automated reasoning with Flyspeck. *J. Autom. Reasoning*, 53(2): 173–213, 2014. doi: 10.1007/s10817-014-9303-3. URL http://dx.doi.org/10.1007/s10817-014-9303-3.

C. Kaliszyk and J. Urban. MizAR 40 for Mizar 40. *J. Autom. Reasoning*, 55(3):245–256, 2015a. doi: 10.1007/s10817-015-9330-8. URL http://dx.doi.org/10.1007/s10817-015-9330-8.

C. Kaliszyk and J. Urban. FEMaLeCoP: Fairly efficient machine learning connection prover. In M. Davis, A. Fehnker, A. McIver, and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*, volume 9450 of *Lecture Notes in Computer Sci-*

ence, pages 88–96. Springer, 2015b. ISBN 978-3-662-48898-0. doi: 10.1007/978-3-662-48899-7_7. URL http://dx.doi.org/10.1007/978-3-662-48899-7_7.

C. Kaliszyk, J. Urban, and J. Vyskocil. Certified connection tableaux proofs for HOL Light and TPTP. In X. Leroy and A. Tiu, editors, *Proceedings of the 2015 Conference on Certified Programs and Proofs, CPP 2015, Mumbai, India, January 15-17, 2015*, pages 59–66. ACM, 2015. ISBN 978-1-4503-3296-5. doi: 10.1145/2676724.2693176. URL http://doi.acm.org/10.1145/2676724.2693176.

L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, editors, *Machine Learning: ECML 2006, 17th European Conference on Machine Learning, Berlin, Germany, September 18-22, 2006, Proceedings*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293. Springer, 2006. ISBN 3-540-45375-X. doi: 10.1007/11871842_29. URL http://dx.doi.org/10.1007/11871842_29.

L. Kovács and A. Voronkov. First-order theorem proving and Vampire. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *LNCS*, pages 1–35. Springer, 2013. ISBN 978-3-642-39798-1.

J. Otten. leanCoP 2.0 and ileanCoP 1.2: High performance lean theorem proving in classical and intuitionistic logic (system descriptions). In A. Armando, P. Baumgartner, and G. Dowek, editors, *Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings*, volume 5195 of *Lecture Notes in Computer Science*, pages 283–291. Springer, 2008. ISBN 978-3-540-71069-1. doi: 10.1007/978-3-540-71070-7_23. URL http://dx.doi.org/10.1007/978-3-540-71070-7_23.

T. Raths and J. Otten. randoCoP: Randomizing the proof search order in the connection calculus. In B. Konev, R. A. Schmidt, and S. Schulz, editors, *Proceedings of the First International Workshop on Practical Aspects of Automated Reasoning, Sydney, Australia, August 10-11, 2008*, volume 373 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2008. URL http://ceur-ws.org/Vol-373/paper-08.pdf.

M. P. D. Schadd, M. H. M. Winands, M. J. W. Tak, and J. W. H. M. Uiterwijk. Single-player Monte-Carlo tree search for SameGame. *Knowl.-Based Syst.*, 34: 3–11, 2012. doi: 10.1016/j.knosys.2011.08.008. URL http://dx.doi.org/10.1016/j.knosys.2011.08.008.

S. Schulz. System description: E 1.8. In K. L. McMillan, A. Middeldorp, and A. Voronkov, editors, *LPAR*, volume 8312 of *LNCS*, pages 735–743. Springer, 2013. ISBN 978-3-642-45220-8. doi: 10.1007/978-3-642-45221-5_49. URL http://dx.doi.org/10.1007/978-3-642-45221-5_49.

D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–503, 2016. URL http://www.nature.com/nature/journal/v529/n7587/full/nature16961.

K. Slind and M. Norrish. A brief overview of HOL4. In O. A. Mohamed, C. Muñoz, and S. Tahar, editors, *TPHOLs 2008*, volume 5170 of *LNCS*, pages 28–32. Springer, 2008.

G. Sutcliffe. The 6th IJCAR automated theorem proving system competition - CASC-J6. *AI Commun.*, 26 (2):211–223, 2013. doi: 10.3233/AIC-130550. URL http://dx.doi.org/10.3233/AIC-130550.

J. Urban, J. Vyskocil, and P. Stepánek. MaLeCoP machine learning connection prover. In K. Brünnler and G. Metcalfe, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 20th International Conference, TABLEAUX 2011, Bern, Switzerland, July 4-8, 2011. Proceedings*, volume 6793 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 2011. ISBN 978-3-642-22118-7. doi: 10.1007/978-3-642-22119-4_21. URL http://dx.doi.org/10.1007/978-3-642-22119-4_21.

M. Wenzel, L. C. Paulson, and T. Nipkow. The Isabelle framework. In O. A. Mohamed, C. A. Muñoz, and S. Tahar, editors, *TPHOLs*, volume 5170 of *LNCS*, pages 33–38. Springer, 2008. ISBN 978-3-540-71065-3.